

AMENDMENTS TO THE CLAIMS

1. (Currently amended): A method in a data processing system that includes a plurality of nodes, each one of said plurality of nodes including a computer and a disk cache coupled to the computer, said system including multiple independent computers that are each coupled to a different disk cache, for managing cached data, the method comprising:

responsive to a first one of said computers in a first node initiating a read operation on a block of data, placing an indication on a directory of data blocks identifying the first one of said computers data processing system as containing a copy of the data block in a disk cache that is coupled to said first one of said computers, said disk cache being included in said first one of said plurality of nodes, including in said indication [[with]] a location in the first one of said computers data processing system in which a flag associated the copy of the data block is located, said flag indicating whether said copy of said data block that is located in said disk cache is valid; [[and]] each one of said computers that is coupled to a disk cache that includes a copy of said data block including a flag that indicates whether said copy of said data block that is stored in the disk cache in each one of said computers is valid; and

responsive to initiating a write operation on the data block, sending a message to all [[other]] others of said computers data processing systems identified in the directory of data blocks as containing a copy of the data block to reset the flag included in said others of said computers such that the flag indicates that the data in the copy of the data block that is stored in a disk cache that is coupled to each one of said others of said computers is invalid without requiring any action by the others of said computers that received other data processing systems receiving the message.
2. (Currently amended): The method of claim 1, wherein the directory of data blocks is located in at least one of the computers data processing system and the other data processing systems.
3. (Original): The method of claim 1, wherein the data block is a page.

4. (Currently amended): A method in a data processing system for managing a plurality of disk caches of data, the system including a plurality of nodes of independent computers, each one of said plurality of nodes including one of said plurality of disk caches, wherein the data processing system includes a plurality of processors, the method comprising:

dedicating one of said independent computers a processor within the plurality of processors for polling for request messages from other computers data processing systems;

responsive to initiating a read operation to read data on a data block into one of said plurality of disk caches, posting an indication on a directory of data blocks identifying the read operation by identifying one of the computers that initiated the read operation the data processing system; and

responsive to initiating a write operation on the data block, sending a message to all of the other computers data processing systems that the data block is invalid.

5. (Currently amended): The method of claim 4 further comprising:

obtaining a lock on a data block by one of said computers included in one of said plurality of nodes;

determining whether a copy of the data block is stored present within a local disk cache included in said one of said plurality of said nodes; and

responsive to a copy of the data block not being stored within absent from the local disk cache, checking a validity of the data block utilizing said directory in the buffer.

6. (Currently amended): The method of claim 4 further comprising:

providing a lock table, wherein the lock table contains data identifying a lock that is held on the data block and includes the data processing system and a location of a validity flag in each computer system in one of said plurality of nodes that includes one of said disk caches in which a copy of said data block is currently stored, each said validity flag in a computer system indicating in which the validity flag indicates whether the data block that is currently stored in said one of said plurality of disk caches is valid.

7. (Original): The method of claim 4, wherein the data block is a page.
8. (Original): The method of claim 4, wherein the message initiates an invalidation of the data block.
9. (Currently amended): A method in a system that includes multiple independent computer systems ~~a data processing system~~ for managing data in ~~a distributed buffer system~~, the method comprising:
 - identifying an operation to access [[the]] data;
 - determining whether a copy of the data is stored locally in a disk cache in a node ~~that includes said disk cache coupled to one of the computers present locally within the data processing system~~;
 - responsive to the copy of the data being stored in said local disk cache present locally within the data processing system, checking an indicator for the data to determine whether the copy of the data is valid; [[and]]
 - responsive to the data being valid, accessing, by the one of the computers, the copy of the data[[.]];
 - responsive to the accessing being a write access, obtaining identification of each one of the computers that is coupled to a disk cache in which a copy of the data is currently stored to form a set of identified computers; and
 - setting indicators in the set of identified computers to indicate that a copy of the data that is stored in said disk cache of each one of the computers is invalid.
10. (Currently amended): The method of claim 9 further comprising:
 - responsive to an absence of a copy of the data being stored in said local disk cache present locally within the data processing system, copying the data into the local disk cache data processing system; and
 - setting one of said indicators ~~an indicator~~ to indicate that the data copied into the local disk cache data processing system is valid.
- 11-13. (Canceled)

14. (Original): The method of claim 9, wherein the data is a page.
15. (Original): The method of claim 9, wherein the operation is read operation.
16. (Original): The method of claim 9, wherein the operation is a write operation.
17. (Currently amended): A data processing system comprising:
~~a bus system;~~
~~a communications unit connected to the bus system;~~
~~a memory connected to the bus system, wherein the memory includes a set of instructions; and~~
~~a plurality of nodes including multiple independent computers, each one of said plurality of nodes including a different disk cache coupled to each one of said computers;~~
~~each one of the computers including a processing unit connected to the bus system, wherein the processing unit executes a [[the]] set of instructions to, place an indication on a directory of data blocks identifying one of the computers that includes said processor the data processing system as containing a copy of a data block in a disk cache that is coupled to said one of said computers, said disk cache being included in one of said plurality of nodes, the indication including [[with]] a location in the one of the computers data processing system in which a flag associated the copy of the data block is located, the flag indicating whether the copy of the data block that is locked in the disk cache in the one of the plurality of nodes is valid in response to initiating a read operation on the block of data; and send a message to all others of the computers other data processing systems identified in the directory of data blocks as containing a copy of the data block to reset the flag included in the others of the computers such that the flag indicates that the data in the copy of the data block that is stored in a disk cache that is coupled to the others of the computers is invalid without requiring any action by the others of the computers that received other data processing systems receiving the message in response to initiating a write operation on the data block.~~

18-19. (Canceled)

20. (Currently amended): A data processing system that includes a plurality of nodes, each one of said plurality of nodes including a computer and a disk cache coupled to said computer, said system including multiple independent computers, for managing cached data, the data processing system comprising:

responsive to a first one of said computers in a first node initiating a read operation on a block of data, placing means for placing an indication on a directory of data blocks identifying the first one of said computers as containing a copy of the data block in a disk cache that is coupled to said first one of said computers, said disk cache being included in said first one of said plurality of nodes, including in said indication a location in the first one of said computers in which a flag associated the copy of the data block is located, said flag indicating whether said copy of said data block that is located in said disk cache is valid;

each one of said computers that is coupled to a disk cache that includes a copy of said data block including a flag that indicates whether said copy of said data block that is stored in the disk cache in each one of said computers is valid; and

responsive to initiating a write operation on the data block, sending means for sending a message to all others of said computers identified in the directory of data blocks as containing a copy of the data block to reset the flag included in said others of said computers such that the flag indicates that the data in the copy of the data block that is stored in a disk cache that is coupled to each one of said others of said computers is invalid without requiring any action by the others of said computers that received the message.

placing means, responsive to initiating a read operation on a block of data, for placing an indication on a directory of data blocks identifying the data processing system as containing a copy of the data block with a location in the data processing system in which a flag associated the data block is located; and

sending means, responsive to initiating a write operation on the data block, for sending a message to all other data processing systems identified in the directory of data blocks as containing a copy of the data block to reset the flag such that the flag indicates that the data in the data block is invalid without requiring any action by the other data processing systems receiving the message.

21. (Currently amended): The ~~data processing system~~ of claim 20, wherein the directory of data blocks is located in at least one of the ~~computers~~ ~~data processing system~~ and the other ~~data processing systems~~.
22. (Currently amended): The ~~data processing system~~ of claim 20, wherein the data block is a page.
23. (Currently amended): A ~~data processing system~~ for managing a plurality of disk caches of data, the system including a plurality of nodes of independent computers, each one of said plurality of nodes including one of said computers and one of said plurality of disk caches coupled to said one of said computers, wherein the ~~data processing system~~ includes a plurality of processors, the ~~data processing system~~ comprising:
 - dedicating means for dedicating one of said independent computers for polling for request messages from other computers;
 - responsive to initiating a read operation to read data on a data block into one of said plurality of disk caches, posting means for posting an indication on a directory of data blocks identifying the read operation by identifying one of the computers that initiated the read operation; and
 - responsive to initiating a write operation on the data block, sending means for sending a message to all of the other computers that the data block is invalid.
 - dedicating means for dedicating a processor within the plurality of processors for polling for request messages from other data processing systems;
 - posting means, responsive to initiating a read operation to read data on a data block, for posting an indication on a directory of data blocks identifying the read operation by the data processing system; and
 - sending means, responsive to initiating a write operation on the data block, for sending a message to all of the other data processing systems that the data block is invalid to remove the data block from the directory of data blocks.

24. (Currently amended): The data processing system of claim 23 further comprising:
obtaining means for obtaining a lock on a data block by one of said computers
including in one of said plurality of nodes;

determining means for determining whether a copy of the data block is stored
present within a local disk cache included in said one of said plurality of said nodes; and
checking means, responsive to a copy of the data block not being stored within
absent from the local disk cache, for checking a validity of the data block utilizing said
directory in the buffer.

25. (Currently amended): The data processing system of claim 23 further comprising:
providing a lock table, wherein the lock table contains data identifying a lock that
is held on the data block and includes a location of a validity flag in each computer
system in one of said plurality of nodes that includes one of said disk caches in which a
copy of said data block is currently stored, each said validity flag in a computer system
indicating whether the data block that is currently stored in said one of said plurality of
disk caches is valid.

providing means for providing a lock table, wherein the lock table contains data
identifying the data processing system and a location of a validity flag in which the
validity flag indicates whether the data block is valid.

26. (Original): The data processing system of claim 23, wherein the data block is a page.

27. (Original): The data processing system of claim 23, wherein the message initiates an invalidation of the data block.

28. (Currently amended): A data processing system that includes multiple
independent computer systems for managing data in a distributed buffer system, the data processing system comprising:
identifying means for identifying an operation to access data;

determining means for determining whether a copy of the data is stored locally in a disk cache in a node that includes said disk cache coupled to one of the computers;

responsive to the copy of the data being stored in said local disk cache, checking means for checking an indicator for the data to determine whether the copy of the data is valid;

responsive to the data being valid, accessing means for accessing, by the one of the computers, the copy of the data;

responsive to the accessing being a write access, obtaining means for obtaining identification of each one of the computers that is coupled to a disk cache in which a copy of the data is currently stored to form a set of identified computers; and

setting means for setting indicators in the set of identified computers to indicate that a copy of the data that is stored in said disk cache of each one of the computers is invalid.

identifying means for identifying an operation to access the data;

determining means for determining whether a copy of the data is present locally within the data processing system;

checking means, responsive to the copy of the data being present locally within the data processing system, for checking an indicator for the data to determine whether the copy of the data is valid; and

accessing means, responsive to the data being valid, for accessing the copy of the data.

29. (Currently amended): The data processing system of claim 28 further comprising:
copying means, responsive to an absence of a copy of the data being stored in a disk cache present locally within the one of the computers data processing system, for copying the data into the local disk cache data processing system; and

setting means for setting one of said indicators an indicator to indicate that the data copied into the local disk cache data processing system is valid.

30-32. (Canceled)

33. (Original): The data processing system of claim 28, wherein the data is a page.

34. (Original): The data processing system of claim 28, wherein the operation is read operation.

35. (Original): The data processing system of claim 28, wherein the operation is a write operation.

36. (Currently amended): A computer program product in a computer readable medium for managing a plurality of caches data, the computer program product comprising:

a system that includes a plurality of nodes, each one of said plurality of nodes including a computer and a disk cache, said system including multiple independent computers;

responsive to a first one of said computers in a first node initiating a read operation on a block of data, instruction means for placing an indication on a directory of data blocks identifying the first one of said computers as containing a copy of the data block in a disk cache that is coupled to said first one of said computers, said disk cache being included in said first one of said plurality of nodes, including in said indication a location in the first one of said computers in which a flag associated the copy of the data block is located, said flag indicating whether said copy of said data block that is located in said disk cache is valid;

each one of said computers that is coupled to a disk cache that includes a copy of said data block including a flag that indicates whether said copy of said data block that is stored in the disk cache in each one of said computers is valid; and

responsive to initiating a write operation on the data block, instructions for sending a message to all others of said computers identified in the directory of data blocks as containing a copy of the data block to reset the flag included in said others of said computers such that the flag indicates that the data in the copy of the data block that is stored in a disk cache that is coupled to each one of said others of said computers is

invalid without requiring any action by the others of said computers that received the message.

~~first instructions, responsive to initiating a read operation on a block of data, for placing an indication on a directory of data blocks identifying the data processing system as containing a copy of the data block with a location in the data processing system in which a flag associated the data block is located; and~~

~~second instructions, responsive to initiating a write operation on the data block, for sending a message to all other data processing systems identified in the directory of data blocks as containing a copy of the data block to reset the flag such that the flag indicates that the data in the data block is invalid without requiring any action by the other data processing systems receiving the message.~~

37. (Currently amended): A computer program product in a computer readable medium for managing a plurality of disk caches of data, wherein the data processing system includes a plurality of processors, the computer program product comprising:

the system including a plurality of nodes of independent computers, each one of said plurality of nodes including one of said plurality of disk caches;

instructions for dedicating one of said independent computers for polling for request messages from other computers;

responsive to initiating a read operation to read data on a data block into one of said plurality of disk caches, instructions for posting an indication on a directory of data blocks identifying the read operation by identifying one of the computers that initiated the read operation; and

responsive to initiating a write operation on the data block, instructions for sending a message to all of the other computers that the data block is invalid.

~~first instructions for dedicating a processor within the plurality of processors for polling for request messages from other data processing systems;~~

~~second instructions, responsive to initiating a read operation to read data on a data block, for posting an indication on a directory of data blocks identifying the read operation by the data processing system; and~~

~~third instructions, responsive to initiating a write operation on the data block, for sending a message to all of the other data processing systems that the data block is invalid to remove the data block from the directory of data blocks.~~

38. (Currently amended): A computer program product in a computer readable medium for managing data ~~in a distributed buffer system~~, the computer program product comprising:

a system that includes multiple independent computer systems;
instructions for identifying an operation to access data;
instructions for determining whether a copy of the data is stored locally in a disk cache in a node that includes said disk cache coupled to one of the computers;
responsive to the copy of the data being stored in said local disk cache,
instructions for checking an indicator for the data to determine whether the copy of the data is valid;
responsive to the data being valid, instructions for accessing, by the one of the computers, the copy of the data;
responsive to the accessing being a write access, instructions for obtaining identification of each one of the computers that is coupled to a disk cache in which a copy of the data is currently stored to form a set of identified computers; and
instructions for setting indicators in the set of identified computers to indicate that a copy of the data that is stored in said disk cache of each one of the computers is invalid,
~~first instructions for identifying an operation to access the data;~~
~~second instructions for determining whether a copy of the data is present locally within the data processing system;~~
~~third instructions, responsive to the copy of the data being present locally within the data processing system, for checking an indicator for the data to determine whether the copy of the data is valid; and~~
~~fourth instructions, responsive to the data being valid, for accessing the copy of the data.~~